

### IN THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

1. (previously presented): A graphics accelerator, comprising:
  - a plurality of specialized processing subunits, interconnected through a serial message-passing interface to provide a generally pipelined graphics accelerator architecture; and
  - a memory interface which provides a high bandwidth interface independent of said serial message-passing interface directly to a memory associated with said graphics accelerator, said memory capable of storing displayable pixel information.
2. (original): The accelerator of Claim 1, wherein ones of said subunits are configured so that said memory interface accesses multiple tiles of pixels simultaneously.

3. (previously presented): A graphics accelerator, comprising:
- a plurality of specialized processing subunits,
    - interconnected through a serial message-passing interface to provide a reconfigurably pipelined graphics accelerator architecture;
    - at least one of said specialized processing subunits comprising multiple subprocessors connected to operate in parallel on separate tasks; and
  - a high bandwidth memory interface independent of said serial message-passing interface which interfaces to a memory of said graphics accelerator, said memory capable of storing displayable pixel information;
- wherein said serial interface also permits downloading of image data to ones of said subunits.
4. (previously presented): The accelerator of Claim 3, wherein ones of said subunits are configured so that said memory interface accesses multiple tiles of pixels simultaneously.
5. (previously presented): The accelerator of Claim 3, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultaneously.
6. (previously presented): The accelerator of Claim 3, wherein said subunits include a current parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.

7. (previously presented): The accelerator of Claim 1, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultaneously.
8. (previously presented): The accelerator of Claim 1, wherein said subunits include a current parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.
9. (previously presented): A graphics rendering system comprising:  
a host processor;  
a system memory; and  
a graphics accelerator comprising:  
a plurality of specialized processing subunits, interconnected through a message-passing interface to provide a generally pipelined graphics accelerator architecture; and  
a memory interface which provides a high bandwidth interface independent of said serial message-passing interface directly to a memory associated with said graphics accelerator, said memory capable of storing displayable pixel information.
10. (previously presented): The system of Claim 9, wherein ones of said subunits are configured so that said memory interface accesses multiple tiles of pixels simultaneously.
11. (previously presented): The system of Claim 9, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultaneously.

12. (previously presented): The system of Claim 9, wherein said subunits include a current parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.
13. (previously presented): A graphics rendering system comprising:  
a host processor;  
a system memory; and  
a graphics accelerator comprising:  
a plurality of specialized processing subunits,  
interconnected through a serial message-passing interface to provide  
a reconfigurably pipelined graphics accelerator architecture;  
at least one of said specialized processing subunits comprising  
multiple subprocessors connected to operate in parallel on  
separate tasks; and  
a high bandwidth memory interface independent of said serial  
message-passing interface which interfaces to a memory of  
said graphics accelerator, said memory capable of storing  
displayable pixel information;  
wherein said serial interface also permits downloading of image data  
to ones of said subunits.
14. (previously presented): The system of Claim 13, wherein ones of said  
subunits are configured so that said memory interface accesses multiple  
tiles of pixels simultaneously.

15. (previously presented): The system of Claim 13, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultaneously.
16. (previously presented): The system of Claim 13, wherein said subunits include a current parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.
17. (previously presented): A computing system for graphics rendering, the system comprising:  
a host processor;  
a system memory; and  
a computer graphics pipeline coupled to said host processor;  
the computer graphics pipeline comprising:  
a plurality of specialized processing subunits, interconnected through a serial message-passing interface to provide a generally pipelined graphics accelerator architecture; and  
a memory interface which provides a high bandwidth interface independent of said serial message-passing interface directly to a memory associated with said graphics accelerator, said memory capable of storing displayable pixel information.
18. (previously presented): The system of Claim 17, wherein ones of said subunits are configured so that said memory interface accesses multiple tiles of pixels simultaneously.

19. (previously presented): The system of Claim 17, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultancously.
20. (previously presented): The system of Claim 17, wherein said subunits include a curmet parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.
21. (previously presented): A computing system for graphics rendering, the system comprising:  
a host processor;  
a system memory; and  
a computer graphics pipeline coupled to said host processor;  
the computer graphics pipeline comprising:  
a plurality of specialized processing subunits,  
interconnected through a serial message-passing interface to provide a reconfigurably pipelined graphics accclerator architecture;  
at least one of said specialized processing subunits comprising multiple subprocessors connected to operate in parallel on separate tasks; and  
a high bandwidth memory interface independent of said serial message-passing interface which interfaces to a memory of said graphics accelerator, said memory capable of storing displayable pixel information;  
wherein said serial interface also permits downloading of image data to ones of said subunits.

22. (previously presented): The system of Claim 21, wherein ones of said subunits are configured so that said memory interface accesses multiple tiles of pixels simultaneously.
23. (previously presented): The system of Claim 21, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultaneously.
24. (previously presented): The system of Claim 21, wherein said subunits include a current parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.
25. (previously presented): A method of designing graphics accelerators, comprising the steps of:  
interconnecting a plurality of specialized processing subunits through a serial message-passing interface to provide a generally pipelined graphics accelerator; and  
providing a high bandwidth interface independent of said serial message-passing interface directly to a memory associated with said graphics accelerator through a memory interface, said memory capable of storing displayable pixel information.
26. (previously presented): The method of Claim 25, wherein ones of said subunits are configured so that said memory interface accesses multiple tiles of pixels simultaneously.

27. (previously presented): The method of Claim 25, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultaneously.
28. (previously presented): The method of Claim 25, wherein said subunits include a current parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.
29. (**currently amended**): A method of designing graphics accelerators, comprising the steps of:
- interconnecting a plurality of specialized processing subunits through a serial message-passing interface to provide a reconfigurably pipelined graphics accelerator;
  - at least one of said specialized processing subunits comprising multiple subprocessors connected to operate in parallel on separate tasks; and
  - providing ~~[[a]]~~ an interface independent of said serial message-passing interface directly to a memory associated with said graphics accelerator through a memory interface, said memory capable of storing displayable pixel information;
- wherein said serial interface also permits downloading of image data to ones of said subunits.
30. (previously presented): The method of Claim 29, wherein ones of said subunits are configured so that said memory interface accesses multiple tiles of pixels simultaneously.



31. (previously presented): The method of Claim 29, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultaneously.
32. (previously presented): The method of Claim 29, wherein said subunits include a current parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.
33. (previously presented): A graphics rendering method, comprising the steps of:  
receiving graphics primitives; and  
sending said graphics primitives through a graphics accelerator comprising:  
a plurality of specialized processing subunits, interconnected through a serial message-passing interface to provide a generally pipelined architecture; and  
a memory interface which provides a high bandwidth interface independent of said serial message-passing interface directly to a memory associated with said graphics accelerator, said memory capable of storing displayable pixel information.
34. (previously presented): The method of Claim 33, wherein ones of said subunits are configured so that said memory interface accesses multiple tiles of pixels simultaneously.

35. (previously presented): The method of Claim 33, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultaneously.
36. (previously presented): The method of Claim 33, wherein said subunits include a current parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.
37. (previously presented): A graphics rendering method, comprising the steps of:
- receiving graphics primitives; and
  - sending said graphics primitives through a graphics accelerator comprising:
    - a plurality of specialized processing subunits,
    - interconnected through a serial message-passing interface to provide a reconfigurably pipelined graphics accelerator architecture;
    - at least one of said specialized processing subunits comprising multiple subprocessors connected to operate in parallel on separate tasks; and
    - a high bandwidth memory interface independent of said serial message-passing interface which interfaces directly to a memory of said graphics accelerator, said memory capable of storing displayable pixel information;
- wherein said serial interface also permits downloading of image data to ones of said subunits.

38. (previously presented): The method of Claim 37, wherein ones of said subunits are configured so that said memory interface accesses multiple tiles of pixels simultaneously.
39. (previously presented): The method of Claim 37, wherein ones of said subunits are configured so that said memory interface accesses multiple byte-sized tiles of pixels simultaneously.
40. (previously presented): The method of Claim 37, wherein said subunits include a current parameter unit, a vertex shading unit, a vertex machine unit, a cull unit, and a geometry unit.